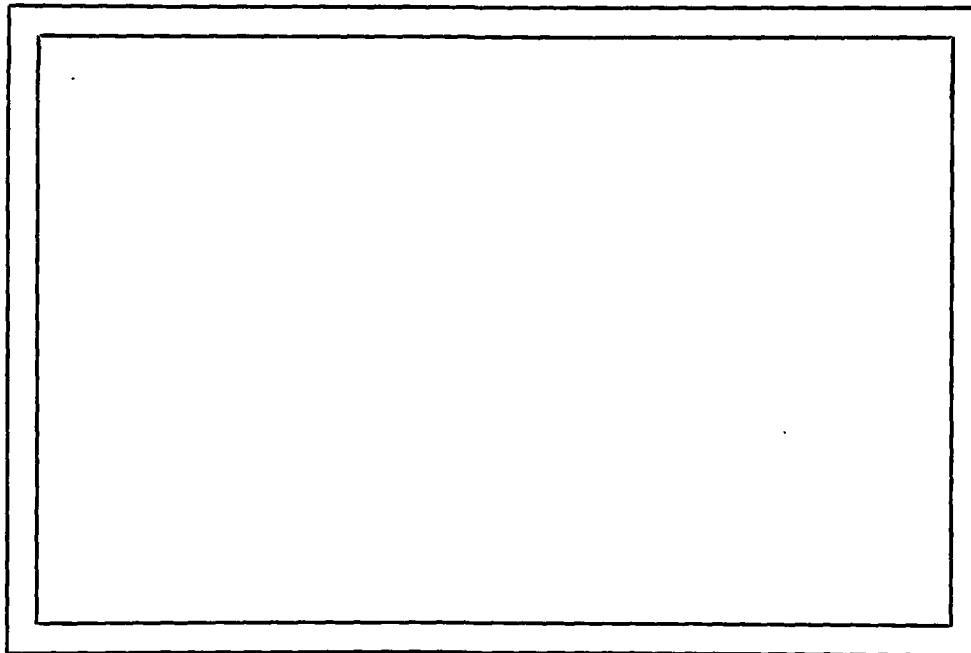


LEVEL *it*

(1)

ADA 086191



COMPUTER SCIENCE  
TECHNICAL REPORT SERIES



UNIVERSITY OF MARYLAND  
COLLEGE PARK, MARYLAND

20742

DTIC  
ELECTE

MAR 3 1980

A

DDC FILE COPY

DISTRIBUTION STATEMENT A

Approved for public release;  
Distribution Unlimited

80 6 30 085

TR-808  
DAAG-53-76C-0138

September 1979

USING PYRAMIDS TO DEFINE  
LOCAL THRESHOLDS FOR BLOB DETECTION

Michael Shneier  
Computer Vision Laboratory  
Computer Science Center  
University of Maryland  
College Park, MD 20742

*See 1473 in back*

ABSTRACT

A method for detecting blobs in images is described. The method involves building a succession of lower resolution images and looking for spots in these images. A spot in a low-resolution image corresponds to a distinguished compact region in a known position in the original image. Further, it is possible to calculate thresholds in the low-resolution image, using very simple methods, and to apply those thresholds to the region of the original image corresponding to the spot. Examples are shown in which the technique is applied to several images.

DTIC  
S ELECTE D  
JUL 3 1980  
A

The support of the Defense Advanced Research Projects Agency and the U.S. Army Night Vision Laboratory under Contract DAAG-53-76C-0138 (DARPA Order 3206) is gratefully acknowledged, as is the help of Kathryn Riley in preparing this paper.

DISTRIBUTION STATEMENT A

Approved for public release;  
Distribution Unlimited

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DDC TAB	
Unannounced	
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or special
A	

## 1. Introduction

The most common way to extract objects from a picture is to threshold the picture. Many different techniques have been used to select good thresholds for this purpose [4]. Threshold selection involves choosing a gray level  $t$  such that all gray levels greater than  $t$  are mapped into the "object" label, while all other gray levels are mapped into the "background" label. In its simplest form, a single threshold is chosen for the whole image. This does not usually give good results because of variations in lighting, or because there are several objects in the picture with different gray-level characteristics. For better results, several local thresholds can be extracted from various parts of the picture, and can be applied just in those regions.

This paper describes a method of identifying parts of a picture on which to apply a threshold, and a means of calculating a local threshold for each of these parts. The method involves constructing a "pyramid" of images, each of lower resolution than its predecessor [1-3]. At some level of the pyramid, it is to be expected that any blob-like object should become spot-like. Thus, by running a spot-detector over the low-resolution images, the interesting regions in the picture can be discovered, and only these regions need be thresholded. In addition, the characteristics of the local regions (or the

spots) can be used to calculate a good local threshold.

Examples are given of the application of the method to several images. In all cases the results are quite good, and highlight the usefulness of the method.

## 2. The algorithm

The algorithm has two main tasks. The first is to find parts of the picture that differ significantly from the background (likely objects), while the second is to calculate a local threshold for each of these parts and apply it in the neighborhood of the parts. Both tasks make use of the pyramid of low-resolution images.

1. If the whole pyramid has been constructed, stop.  
Otherwise, read in the previous pyramid level (the picture, if this is the first iteration).
2. Build a new level (see below).
3. Apply a spot detector to the new level.
4. Evaluate the spots resulting from step 3 and find "good" spots (see below). If there are too many good spots, go to 1.
5. For each good spot,
  - a. calculate a threshold (see below);
  - b. apply the threshold to the region in the original picture corresponding to the spot and write the results to the output picture.
6. Go to 1.

The original image forms the base of the pyramid. Each level is constructed on top of its predecessor, and is processed before its successor is constructed. This means that

only one level need be maintained at any time, in addition to the original picture and the partially-constructed thresholded picture.

A pyramid level is constructed from its predecessor by mapping 2 by 2 squares of pixels from the previous level into single pixels in the new level. Two methods of calculating the new value from the old were implemented. The first involves simple averaging of the four pixels. In the second method, each 2 by 2 block of pixels is examined and the four gray levels are sorted in order of brightness. The middle two values are then averaged to give the new pixel corresponding to the 2 by 2 block. This process gives results that maintain edges reasonably well. In practice, both methods usually produce the same results. The new level of the pyramid is one quarter the size of the old (Figure 1a).

Having built a level of the pyramid, the next step is to apply a spot detector to it. The spot detector is a simple mask (Figure 2) that is applied at every point in the image. It looks for points that differ from their neighbors and scores them according to how much they differ. Note that the central value in the mask is smaller than an unbiased mask would require. This is to insure that the spots are more than marginally different from their neighbors. It tends to ignore spots caused by noise. The result of running the spot detector

is a new image with high values where there are spots, and low values elsewhere.

The spot detector is very conservative, so another process is run to find a subset of "good" spots. Good spots are spots that are isolated. At low levels of the pyramid (high resolution), spots that are close together are deleted because they can be expected to merge into single spots higher up in the pyramid. At higher levels of the pyramid, this is not such a good idea because single spots represent large regions in the original picture. Thus, the definition of "good" is weighted by the level in the pyramid. A spot is good if the number of its neighbors that also responded positively to the spot detector is less than a level-dependent threshold.

Each spot in the low-resolution image corresponds to a region in the picture. If there are too many spots, then large parts of the picture will be covered. If there is indeed an object in the picture, it should coalesce into a smaller number of spots higher in the pyramid. If there is no object, then all the spots represent noise. In either case, the picture is too "busy". A maximum number of good spots is allowed at each level. If this number is exceeded, no further processing is performed, and a new pyramid level is constructed.

When a small enough number of good spots is discovered at a given level in the pyramid, the thresholding can be performed. Notice that it need only be applied to the regions in the

picture corresponding to the spots in the pyramid. All other regions are ignored.

Many threshold selection techniques are applicable at this stage. There are the standard techniques [4] which may be applied to the picture itself in the region corresponding to a spot. In addition, it is possible to make use of the information in the low-resolution image to calculate a threshold. Both approaches were followed for the examples to be discussed here. Using the low resolution image has the advantage that simple operations on the low resolution image correspond to complex operations involving much larger numbers of points in the picture.

The simplest threshold that can be extracted from the low resolution image is simply the gray level of the spot. This threshold is equivalent to the average gray level of the region in the picture corresponding to the spot. Usually, this threshold does not extract the whole object because the high gray levels bias the threshold, and there are very few non-object points in the region to provide an opposite bias (Figure 1c).

An alternative threshold is obtained by ignoring the spot itself, and averaging the surrounding points in the low-resolution picture. This suffers from the opposite problem from the previous method. Now, too many non-object points reduce the

threshold, and so parts of the background are classified as belonging to the object (Figure 1d).

A compromise between these two methods gives very good results. The outputs from the above two threshold selection processes are averaged, and the result is used as the threshold (Figure 1b).

The threshold is applied to a region slightly larger than that corresponding to the spot. This is to insure that parts of the object that were averaged into different points in the low-resolution image still may be classified, provided that they are not too far away from the spot center. If, indeed, the object extends a significant distance from the spot center, the spot detector should have found several spots in the neighborhood, each of which would be processed separately (or they would all be merged into a larger spot at the next level).

Another method of calculating a local threshold was also implemented. The method involves computing a histogram of the gray levels in the regions of the original picture that correspond to spots. For each spot a histogram is constructed for a region slightly larger than the projection of the spot onto the picture. The histogram is then examined, and a threshold is selected. The process of selection is complicated by the shape of the histogram, which tends either to be uni-modal, or to have no significant peaks (Figure 3). The method

that was used to find a threshold involves making an initial estimate, and refining the estimate on the basis of the shape of a part of the histogram.

The initial guess that was used was one of the naive thresholds mentioned above. The gray level corresponding to the spot in the pyramid provides an estimate of the gray level in the center of the object. Usually, the estimate needs to be modified to take account of parts of the object close to the background. To accomplish this, the histogram is examined, starting at the initial estimate, and moving in the direction of the background gray levels. The highest peak in the histogram in this direction is discovered, and the final threshold is chosen at the deepest valley between this peak and the initial estimate. This usually results in a good threshold, in most cases in one very similar to the averaging of the center and surround points in the pyramid discussed above.

The output picture is initially blank. The only regions of the picture that are changed are those that correspond to positive responses to the spot detector at some level in the pyramid. As a result, very little background noise appears in the output.

### 3. Examples

The method was applied to 24 FLIR images and to a picture of part of a handwritten signature. The results are shown in Figures 4-7. The examples are divided into three categories.

The first set of pictures (Figure 4) was processed using a simple averaging scheme for building the pyramids. The threshold was selected from the low resolution image by taking the average of the center (spot) gray level, and the average surrounding gray level.

Sometimes, when the contrast between the object and the background is small, the averaging process may cause the object to merge into the background. For FLIR imagery, it was found that it is often better to use the median instead of the average in building the pyramids. Figure 5 shows a set of examples where this was done. The threshold selection used the same method as for Figure 4.

The alternative method of selecting a threshold by examining the histogram is illustrated in Figures 6 and 7. Figure 6 shows four FLIR images and the results of thresholding them. The pyramids for these images were constructed by averaging, and the thresholds were selected by examining a histogram of a region in the image slightly larger than that corresponding to the spot.

Figure 7 illustrates the difference between selecting the threshold using only the low-resolution image, and making use

of the histogram as well. For the signature in Figure 7, the histogram method results in a much cleaner thresholded image.

#### 4. Discussion

The blob-detection system described here is the first stage in a more ambitious feature-detection scheme. As it stands, the system provides a good threshold selection technique, with several advantages. One of the most important advantages is the ability of the system to isolate significant regions in a picture. This results both in better local threshold selection and in cleaner thresholded images. The thresholds are tailored specifically to the region to which they are applied, and uninteresting regions are ignored.

A problem that arose from the way the algorithm was implemented concerns the treatment of points on the borders of the picture. These points were ignored in the implementation, and, as a result, the algorithm discovered significant objects only if they were not on the border of the picture. This effect could be aggravated by the pyramid-building process because a point on the border of an image high in the pyramid corresponds to a fairly large region in the picture. There are several ways of overcoming this problem. For example, one-sided spot detectors could be used at the edges of the pictures, or the pictures could be extended either by reflection about the edge, or by folding the edges over so that the left and right and the top and bottom edges are contiguous.

A question that arises naturally concerns the amount of averaging between levels in the pyramid. Perhaps the exponential

tapering used in these experiments is too harsh, and spot detectors of various intermediate sizes should be used in addition to those used here. This would more accurately capture the fine detail of the shapes and allow greater control over threshold selection. It is expected that further research will be conducted on this aspect of the algorithm.

An extension of the method that is currently under investigation is the detection of elongated objects. In conventional thresholding schemes, the shape of an object can only be discovered after the object has been extracted. It is not possible to search for objects with specific shape properties. Using the current method, however, it is possible to extract only those features that are of the desired shape. For example, to extract elongated objects, a line or streak detector can be applied instead of a spot detector. Preliminary results suggest that a straightforward extension of the blob-detection system can be produced which will detect only the elongated objects in a picture. This will help to alleviate a problem that sometimes arises when objects are not sufficiently blob-like. In such cases, some parts of the object may not be covered by the projection of a spot, and only part of the object may be thresholded.

Eventually, the system is envisaged as having multiple cooperating parts. Several feature detectors will be run at

each level of the pyramid, for example, both line detectors and spot detectors. These would then interact within the levels and across levels. The whole system should be able to detect many different features simultaneously, and classify them on the basis of both local and global information.

## 5. Conclusions

A new method of detecting blobs in a picture by spot detection and local thresholding has been presented. The examples showed how simple threshold-detection calculations on low-resolution images can lead to good segmentation of the picture.

The method readily lends itself to extensions to more complex feature detection tasks, including detection of objects with specific properties, e.g. elongated objects.

It is expected that the method will eventually be included in a comprehensive, multilevel feature-extraction system that makes use of multiple-resolution images and responses from several different feature detectors.

### References

1. A. R. Hanson and E. M. Riseman, Segmentation of natural scenes. In Computer Vision Systems (Hanson and Riseman, eds.), Academic Press, New York, pp. 129-163, 1978.
2. S. L. Tanimoto, Regular hierarchical image and processing structures in machine vision. Ibid., pp. 165-176.
3. L. Uhr, "Recognition cones," and some test results; the imminent arrival of well-structured parallel-serial computers; positions, and positions on positions. Ibid., pp. 363-378.
4. J. S. Weszka, A survey of threshold selection techniques. Computer Graphics Image Processing, 7, pp. 259-265, 1978.



Figure 1

- a) A FLIR image of a tank, and the pyramid constructed from it. b) Thresholded image using the average of center and surrounding spots. c) Thresholded image using surrounding spots only. d) Thresholded image using center spot only.

<u>+1</u>	<u>+1</u>	<u>+1</u>
<u>+1</u>	<u>7</u>	<u>+1</u>
<u>+1</u>	<u>+1</u>	<u>+1</u>

Figure 2

The mask used for the spot detector.

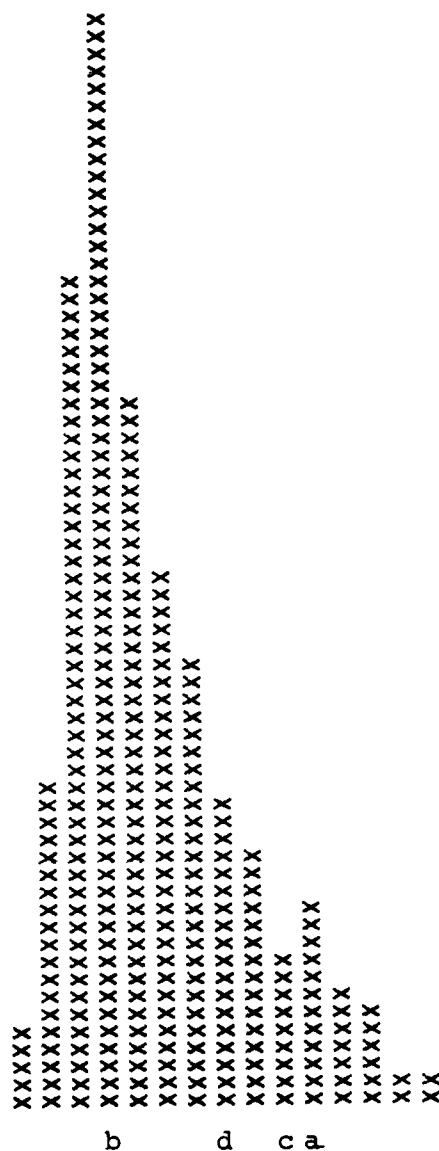


Figure 3

An example of a histogram used for threshold selection. Point a is the initial point chosen for thresholding (see text). Point b is the highest peak in the direction of the background. Point c is the point chosen as the final threshold. Point d is the threshold chosen by the method of averaging the center and background points in the low-resolution image. The histogram is for a spot in the bottom left picture of Figure 6. The small size of the spot results in a very low peak in the histogram (at a).

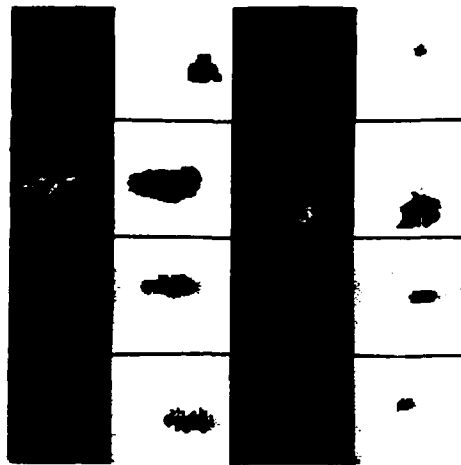


Figure 4

Eight FLIR images and their thresholded outputs. The pyramid was built by averaging in these examples and the threshold was selected as the average of the center and surrounding points in the low-resolution image.

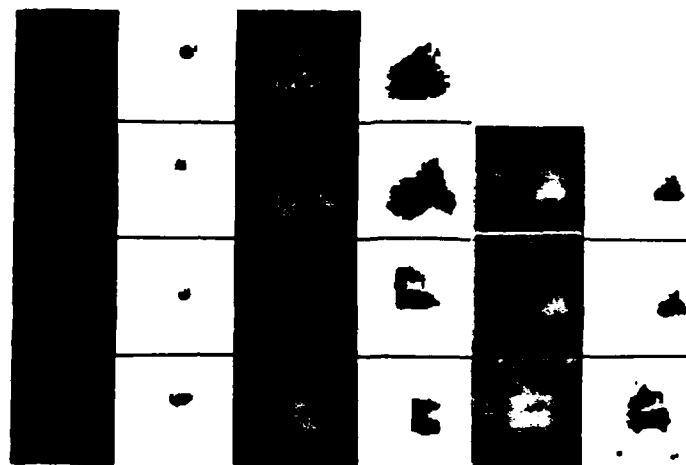


Figure 5

Eleven FLIR images and their thresholded outputs. The pyramid was built using the median and the threshold was selected as the average of the center and surrounding points in the low-resolution image.



Figure 6

Four FLIR images and their thresholded outputs. The pyramids in these examples were constructed by averaging, and the threshold was selected by examining the histograms of local regions corresponding to spots.



a



b



c

Figure 7

- a) A picture of part of a handwritten signature.
- b) The thresholded output using the average of the center and surrounding low-resolution points.
- c) The result of calculating a threshold by examining the histogram.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
	AD-A086794	
4. TITLE (and Subtitle)	5. TYPE OF REPORT & PERIOD COVERED	
(6) USING PYRAMIDS TO DEFINE LOCAL THRESHOLDS FOR BLOB DETECTION.	(9) Technical rept.	
7. AUTHOR(s)	6. PERFORMING ORG. REPORT NUMBER	
(10) Michael/Shneier	(14) TR-808	
(11) Sep 79	8. CONTRACT OR GRANT NUMBER(s)	
	DAAG-53-76C-0138	
9. PERFORMING ORGANIZATION NAME AND ADDRESS	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
Computer Vision Laboratory University of Maryland College Park, MD 20742		
11. CONTROLLING OFFICE NAME AND ADDRESS	12. REPORT DATE	
U.S. Army Night Vision Laboratory Fort Belvoir, VA 22060	September 1979	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	13. NUMBER OF PAGES	
	20	
15. SECURITY CLASS. (of this report)	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
	Unclassified	
16. DISTRIBUTION STATEMENT (of this Report)		
Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
(15) DAAG 53-76-C-0138 ✓ DARPA Order-3206		
18. SUPPLEMENTARY NOTES		
(12) 22		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
Image processing Pattern recognition Segmentation Thresholding Spot detection		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)		
A method for detecting blobs in images is described. The method involves building a succession of lower resolution images and looking for spots in these images. A spot in a low-resolution image corresponds to a distinguished compact region in a known position in the original image. Further, it is possible to calculate thresholds in the low-resolution image, using very simple methods, and to apply those thresholds to the region of the original image corresponding to the spot. Examples are shown in which the technique is applied to several images.		

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

411074

xlt